

Market Segmentation for Privacy Differentiated “Free” Services

Chong Huang, Lalitha Sankar

Abstract

The emerging marketplace for online free services in which service providers earn revenue from using consumer data in direct and indirect ways has lead to significant privacy concerns. This leads to the following question: can the online marketplace sustain multiple service providers (SPs) that offer privacy-differentiated free services? This paper studies the problem of market segmentation for the free online services market by augmenting the classical Hotelling model for market segmentation analysis to include the fact that for the free services market, a consumer values service not in monetized terms but by its quality of service (QoS) and that the differentiator of services is not product price but the privacy risk advertised by a SP. Building upon the Hotelling model, this paper presents a parametrized model for SP profit and consumer valuation of service for both the two- and multi-SP problems to show that: (i) when consumers place a high value on privacy, it leads to a lower use of private data by SPs (i.e., their advertised privacy risk reduces), and thus, SPs compete on the QoS; (ii) SPs that are capable of differentiating on services that do not directly target consumers gain larger market share; and (iii) a higher valuation of privacy by consumers forces SPs with smaller untargeted revenue to offer lower privacy risk to attract more consumers. The work also illustrates the market segmentation problem for more than two SPs and highlights the instability of such markets.

Keywords-free online services, privacy-differentiated services, quality of service, market segmentation.

I. INTRODUCTION

There has been a steady increase in online interactions between consumers and retailers, where the term retailers refers to entities who sell or offer (for free) a product. In fact, many oft used online services are offered for free and consumers implicitly accede to tracking for customized services. Targeted ads are a part of the emerging revenue/profit model for such retailers, henceforth referred to as *service providers* (SPs), especially those offering free services. Consumers are delighted by free services until they begin encountering privacy violations on a daily/frequent basis. While such infractions taken individually could be ignored or discounted, the totality of data available about consumers with a variety of retailers and the resulting privacy consequences raise serious concerns.

Service providers are beginning to acknowledge that consumers are sensitive to privacy violations. For example, Google [1] and Apple [2] recently adopted differentially private mechanisms for collecting user data for statistical analyses. However, the details of these mechanisms are opaque and offer even less clarity on whether the consumer

C. Huang and L. Sankar are with the Department of Electrical, Computer, and Energy Engineering at Arizona State University, Tempe, AZ 85287 (e-mail: chong.huang@asu.edu, lalithasankar@asu.edu).

actually has a choice. In this context, it is worth understanding if privacy-differentiated services can provide such choices for consumers. In a competitive marketplace, the aggregate weight of targeting may drive some customers to seek a more privacy-protective alternative. The cost to the consumer for this action may be a lower quality of service (QoS) (e.g., poorer search engine capabilities). However, it could eventually lead to a more open model for consumer sharing of private information, i.e., one from implicit assent to informed consent.

In this paper, we identify and formalize conditions under which privacy-differentiated services are sustainable and examine when *competition* based on privacy protections could lead to a sustainable marketplace for online free services. As a hypothetical running example we posit a marketplace with two search engines Google Search and DuckDuckGo. The former tracks searches by users thereby offering higher quality service (QoS) (e.g., search accuracy) while the latter explicitly does no tracking and may offer a relatively lower QoS. More generally, our model allows differentiating SPs by their QoS and privacy-sensitive offerings to determine the existence of a stable market for privacy-differentiated services. We apply and build upon classical game theoretic methods, in particular the Hotelling model for market segmentation, to quantify the market segmentation. We also generalize the model to multiple SPs (e.g., Google, DuckDuckGo, and Bing) and illustrate the instability of multi-competitor markets.

A. Related Work

An extensive body of literature on economic models for privacy was recently reviewed by Acquisti *et al.* [3]. These models illustrate the large semantic range covered by the word “privacy”. Targeted advertising is a common method for service providers to exploit knowledge of a consumer in a way that can cause privacy violations. Our work is informed by the literature on targeting strategies for retailers [4]–[13], but rather than optimizing retailer strategies we are interested in identifying how privacy-differentiated services can address privacy concerns.

The problem of market segmentation is a classic and well-studied problem in microeconomics [14] with focus on how pricing and product differentiation can lead to a stable and competitive marketplace. However, the free online services market present a new challenge wherein monetary quantification of both ‘free’ services and the data collected about consumers is not simple and straightforward. Equally challenging is the quantification of consumer privacy since it requires capturing the heterogeneous expressions of privacy sensitivity that can range from ‘don’t care’ at one extreme to ‘hyper vigilant’ at the other. However, some aspects of market models can be brought to bear to our problem; in particular, the oligopolistic market model with a small number of competitors, barriers to entry that are not as high as those for monopolies, and with differentiated products fits appropriately for the markets we are considering wherein two or (a few) more service providers offer products of the same type but differentiated by QoS and privacy risk.

For a two-player oligopolistic market game, the Cournot-Nash and Bertrand duopoly models are considered classic models wherein the two firms differentiate using quality and price, respectively. A more nuanced model that captures differentiation between two firms and consumer preferences is the Hotelling model [15]. This model captures differentiation between market players by mapping firms to positions on a unit length line such that the location is indicative of the firm’s ‘differentiation level’, the total line length is reflective of the entire market, a consumer’s privacy preference is a point on the line, and the optimal locations of the firm resulting from the

simultaneous game between the players indicate the resulting segmentation. The model captures utility for consumer as both the advantage (price, quality, etc) from the firm as well as the ‘distance cost’ from the consumer’s location to that firm. Consumers choose the seller which give them the highest utility (in terms of their valuation of the product and its price as well as the ‘transportation’ costs). The Hotelling model has also been extended to include gradations in product (quality) and customer types via a vertical variant of the model [16] .

Privacy and market segmentation. Jentzsch et.al. [10] propose a model to study competitions between two service providers by taking consumer’s privacy preference (binary choices: low privacy/high privacy) into account using a vertical Hotelling model. Thus, consumers select the service provider based on their privacy concerns and the amount of payment to the service provider. They provide analysis of equilibrium strategies for SPs. In [11], Lee *et al.* study the influence of privacy protection on the segmentation of a duopoly. In their model, firms may offer standard and personalized products with personalized prices to three different types of privacy-sensitive consumers (the ‘unconcerned’ who always share information, ‘pragmatic’ ones who only share if a firm adopts privacy protection, and the ‘fundamentalists’ who never share data). They show that a privacy-friendly firm can enlarge market share by attracting more pragmatists to share personal information. From this expansion it can earn more profits rather than compete with its rival for the other consumers. In contrast to both above-mentioned models, our model differs in focusing on ‘free’ services, and thus, introduces new models for quantifying QoS- and privacy-based differentiators; furthermore, our model generalizes the discrete set of privacy sensitive consumers in [11] to a continuous set of privacy risks thus allowing analysis of over an entire range of privacy expression and present a more nuanced view of how SPs should offer services to all types of consumers.

B. Our Contributions

Our work introduces a game-theoretic interaction model for free online services offered by two or more SPs with the goal of understanding whether privacy-differentiated service offerings have the ability to capture market share. Our model captures a variety of free online services such as search engines, social networking sites, and software apps that are free, and therefore, use consumer data in a variety of ways for revenue generation. Specifically, our model is based on the ‘spatial’ Hotelling model wherein the location is now proxy for both the privacy risk levels that the SPs offer and consumers prefer (both often at odds). The QoS of the service now models the classical product price. Our model differentiates itself from the Hotelling model in the following sense: unlike the classical model of non-negative transportation costs from consumer ‘location’ (preference) to either SP ‘location’, a consumer with a specific privacy risk choice gains from choosing an SP with a lower risk offering and loses from choosing one with a higher risk offering. This in turn leads to different outcomes than the classical model; we use a three-stage sequential game to compute the optimal strategies for the SPs and the resulting market share for specific models of cost and revenue (to SPs), distribution of consumer heterogeneous privacy choices, as well as QoS valuation (to consumers).

We present closed form solutions for the two SP market with linear valuation functions (cost, revenue, consumer utility) and a uniform distribution of consumer preferences; for this settings, our results highlight the following: (i) when consumer place a high value on privacy, it leads to a lower use of private data by SPs, i.e., their advertised

privacy risk reduces; (ii) SPs offering high privacy risk services are sustainable only if they offer sufficiently high QoS; (iii) SPs that are capable of differentiating on services that do not directly use consumer data gain larger market share; and (iv) higher consumer privacy valuation forces SPs with smaller privacy-independent (untargeted) revenue to offer lower privacy risk service to attract more consumers. In extending the work to more than two SPs, we illustrate the instability of such markets and highlight the challenges of studying market segmentation for more than two participants (a problem acknowledged in economics [17]).

C. Organization of the Paper

The paper is organized as follows: Section II introduces the system model and the non-cooperative game formulation. The main result for a two-SP market with linear valuation functions are presented in Section III. Section IV discuss equilibrium results for a market with multiple SPs. Finally, concluding remarks and future work are provided in Section V.

II. PROBLEM MODEL AND GAME FORMULATION

Our parametrized model detailed below captures the following privacy-differentiated market segmentation problem: service providers offer free services differentiated by QoS and privacy risks. Online services that are offered for free often generate revenue by using the data they obtain from their consumers. Their gain from using consumer data is captured by a revenue function and their cost of doing so is captured by a cost function. The goal of each SP is to choose a QoS and privacy risk tuple that maximizes its profit (difference of revenue and cost). The heterogeneous expression of consumer privacy sensitivity is modeled as a (probability) distribution of the population over a range of privacy risk values. The consumer will choose the SP that maximizes a desired function of QoS, the privacy risk tolerance of the consumer, and the privacy risk offered by the SP. A stable strategy i.e., a strategy from which no participant will deviate without reduction in utility, of such a non-cooperative game will yield an optimal partition (market segmentation) of the consumers. We build upon the classical Hotelling model proposed by Hotelling in [15] to study market segmentation.

Formally, we introduce a game-theoretic model for two SPs and infinitely many consumers. Each SP offers the same type of free services (e.g., search engine, social network) with a certain privacy risk guarantee ε and quality of service (QoS) v . Thus, an SP differentiates its service by a tuple (ε, v) that it advertises to all consumers. The goal of this work is to determine the fraction of consumers (market segment) that choose each SP when the SPs offer (ε, v) tuples that maximize their profit.

A. Two-SP Market Model

1) *SP Model:* We consider two rational (i.e., profit maximization entities) SPs, denoted by SP_1 and SP_2 . Both SPs provide the same kind of free service; but they differ in the QoS offered. Thus, SP_1 and SP_2 offer QoS v_1 and v_2 , respectively, where in general $v_1 \neq v_2$. Furthermore, SP_1 and SP_2 guarantee that the privacy risk for using their services is at most ε_1 and ε_2 , respectively, where $\varepsilon_1, \varepsilon_2 \in [0, \bar{\varepsilon}]$. Without loss of generality, we assume $\varepsilon_2 \geq \varepsilon_1$. Under this assumption, SP_2 must offer a higher QoS ($v_2 \geq v_1$). Otherwise, its strategy will be dominated by

its opponent since SP_1 will offer both higher QoS and lower privacy risk. For example, SP_1 and SP_2 could be Duckduckgo and Google, respectively, in the search engine market, with the QoS given by the accuracy of search results. On the other hand, the privacy risk can correspond to different guarantees they provide on consumer data use; e.g., whether they will use consumer data only for statistical purposes or target consumers with tailored ads. We model this privacy risk guarantee as a variable taking values over a continuous range. In practice, such guarantees may be coarse granular choices; for example, between completely opting out of the targeting or allowing data use only for statistical purposes or complete data use only by SP or all possible data usage and sale. We assume that the SPs generate revenue in two ways: (i) by exploiting the private data of consumers to offer *targeted* ads and other services to consumers; and (ii) by providing interested advertisers an online platform to reach consumers. This latter revenue is independent of private data and simply derived from the revenue capability of the platform.

Let $R_P(\varepsilon_i)$ denote the revenue of SP_i , $i \in \{1, 2\}$, resulting from using the private data of consumers and let $R_{NP,i}$ denote the revenue generated without using consumers' private information (e.g., from interested advertisers). The total revenue, $R(\varepsilon_i)$, of SP_i from offering privacy guaranteed service is thus

$$R(\varepsilon_i) = R_P(\varepsilon_i) + R_{NP,i}, i \in \{1, 2\}. \quad (1)$$

In reality, offering free services to consumers often comes with a cost to the SPs, such as the cost of service and online platform creation and continued operations. Furthermore, we note that free online services profit from using consumer data and therefore encumber data processing related costs. Let $C(v_i; \varepsilon_i)$ denote the cost of offering free services with privacy risk level ε_i . We model $C(v_i; \varepsilon_i)$ as sum of two non-negative costs: (i) $C_{QoS}(v_i)$ of providing services with QoS v_i ; and (ii) $C_P(\varepsilon_i)$ as the processing (data analytics) cost of exploiting private data to the level of ε_i such that

$$C(v_i; \varepsilon_i) = C_{QoS}(v_i) + C_P(\varepsilon_i), i \in \{1, 2\}. \quad (2)$$

Thus, via (1) and (2), our model captures the fact that the benefit of using private data by each SP involves both cost and revenue.

2) *Consumer Model:* We build upon the classical Hotelling model to formulate both consumer utility and the resulting consumer-SP game. The Hotelling model maps retailers to two locations (x_1, x_2) on a $[0, 1]$ line such that the strategy of each retailer is to determine the best location-price tuple that maximizes its profit. The location (see Figure 1a) is a proxy for a specific product differentiator. A consumer with its own product differentiator preference (traditionally assumed to be uniformly distributed over $[0, 1]$) is mapped to a location $x \in [0, 1]$ on the line as shown in Figure 1a. Such a spatial model allows computing the market segment by identifying both the optimal locations of the retailers and an indifferent threshold between the two optimal retailer locations at which both retailers are equally desirable. For such a uniform consumer preference model, the segmentation for each retailer is simply its distance to the indifference point (see Figure 1a). Consumers choose the retailer with the least product price and “transportation cost” (modeled as a linear function of location) for a desired consumer valuation of the product. Note that transportation costs are metaphorical for any non-price-based differentiation of the two retailers.

For our problem, we obtain a Hotelling model by: (i) introducing a *normalized privacy risk* and mapping it to spatial location; and (ii) by viewing the QoS as the net valuation of service by the consumer. Note that since we study a free services market, we use QoS as a measure of consumer satisfaction. We note that in the classical Hotelling model, the consumer pays a non-negative transportation cost for any retailer whose location is different from its own. However, our problem departs from this model in that higher and lower privacy risks offered by SPs relative to a consumer preferred privacy risk choice are not viewed similarly.

We assume there exists infinitely many rational consumers that are interested in the services provided by the SPs. In keeping the standard game-theoretic definition, rational refers to consumers interested in maximizing some measure of utility via interactions with the SPs. We use a random variable $E \in [0, \bar{\varepsilon}]$ to denote the heterogeneous privacy choices of consumers; such a model assumes that the privacy preferences of consumers are independent and identically distributed, a reasonable assumption when the consumer set is very large. Let $E = \varepsilon$ denote the privacy risk preference of a consumer. If an SP_i offers a privacy risk guarantee ε_i higher than ε , then using its service will result in a privacy cost to the consumer due to perceived privacy risk violation. On the other hand, the consumer gains from choosing an SP_i that offers an $\varepsilon_i < \varepsilon$ as a result of the extra privacy protection offered. Let $x = F_E(\varepsilon) \in [0, 1]$ be a differentiable cumulative distribution function of ε . Thus, $x = F_E(\varepsilon)$ can be considered as a normalized privacy risk tolerance (i.e., restricted to $[0, 1]$) which indicates the proportion of the consumers with a privacy risk tolerance of at most ε . Since the privacy risk ε_i can be over an arbitrary range $[0, \bar{\varepsilon}]$, the normalized spatial privacy risk is given by the CDF $F_E(\varepsilon_i)$. Thus, for a consumer whose normalized privacy risk tolerance is located at $x \in [0, 1]$, its actual privacy risk tolerance is $\varepsilon = F_E^{-1}(x)$. We can similarly map the privacy risks offered by the SPs to *normalized locations* $x_1 = F_E(\varepsilon_1)$ and $x_2 = F_E(\varepsilon_2)$ on the $[0, 1]$ line as shown in Figure 1b.

Analogous to the Hotelling model, we let $u_i(x)$ denote the utility (in units of QoS) from SP_i as perceived by a consumer with a normalized privacy preference (location) x . Our model for $u_i(x)$ contains two parts: (i) a positive QoS v_i offered by SP_i ; and (ii) the gain or loss in the perceived QoS as a result of a mismatch between consumer privacy preference and SP_i 's privacy risk offering. We introduce a gain factor t that allows mapping the privacy mismatch $t(x - x_i)\varepsilon_i$ to a QoS quantity. This mismatch utility indicates that when the SP offers a service with privacy risk lower than the consumer's tolerance, the consumer receives a positive utility due to extra privacy protection. However, if the service offered has a higher privacy risk than the consumer's tolerance, the consumer will receive negative utility for privacy risk violation. In other words, given the same level of QoS, the better the privacy risk guarantee an SP offers, the more the consumer prefers the SP. We now write the utility or profit function for both consumers and SPs.

3) *Consumer utility and SP profits:* For the consumer located at x , the overall perceived utility for choosing services provided by SP_1 and SP_2 are

$$u_i(x) = v_i + t(x - x_i)\varepsilon_i, i \in \{1, 2\}. \quad (3)$$




Figure 1: User choice model for using different SPs

For each $SP_i, i \in \{1, 2\}$, let $(v_{-i}, \varepsilon_{-i})$ be its competitor's strategy. For the revenue and cost models in (1) and (2), the profit of SP_i is simply the difference

$$\pi_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i}) = [R(\varepsilon_i) - C(v_i; \varepsilon_i)]n_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i}), i \in \{1, 2\}, \quad (4)$$

where $n_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i})$ denotes the fraction of consumers who choose SP_i .

Modelling Assumption 1: We assume that the services provided by both SPs have non-negative QoS.

Since consumers are rational, they expect to have positive utility through the interactions with the SPs. It is reasonable to assume that SPs have no incentive to offer services with a negative QoS. In other words, we assume $v_1 \geq 0$ and $v_2 \geq 0$.

Modelling Assumption 2: We assume the model parameters are chosen such that they ensure the market is completely covered by SP_1 and SP_2 .

The above assumption implies that each consumer must choose one of the SPs. Such an assumption is implicitly built into the classical Hotelling model to ensure competition between SPs and our model continues to do so too. Later we provide a sufficient condition for sustaining the equilibrium market segmentation under these assumptions.

B. Two-SP Non-cooperative Game Formulation

We note that the SPs compete against each other through their distinct QoS and privacy risk offerings, which in turn affects consumer choices and helps determine the stable market segmentation. Thus, the interactions between SP_1 and SP_2 can be formulated as a non-cooperative game in which the strategy (action) of each SP is a (privacy risk, QoS) tuple and that of the consumer is choosing an SP. Furthermore, we assume that the SPs are rational and have perfect information, implying that they play to maximize their own profits and know the exact profit function for any given strategy.

The interactions between retailers and consumers in the Hotelling model can be viewed as a sequential game [15]. For our model, such a sequential game involves three stages. In the first stage, the differentiator, i.e. the normalized privacy risk ε_i , is advertised by SP_i . Thus is followed by each SP determining its QoS for the advertised risk. Finally, the consumers choose the preferred SP based on the (ε_i, v_i) tuple that maximizes its utility.

The game can be formally described as follows: (i) a set of players $\{1, 2, \mathcal{C}\}$, where 1 and 2 denote SP_1 and SP_2 , respectively, and the set \mathcal{C} contains infinitely many consumers; (ii) a collection of strategy tuples $(\varepsilon_i, v_i) \in \mathcal{E}_i \times \mathcal{V}_i$

for SP_i and a collection of binary choices (strategies) for the consumer $b \in \mathcal{B} = \{1, 2\}$; and (iii) a profit function π_i for each SP_i and a utility function u_i for each consumer for choosing SP_i .

C. The Subgame Perfect Nash Equilibrium for the Two-SP Game

In a sequential game, each stage is referred to as a subgame [18]. One often associates a strategy profile with a sequential game. A strategy profile is a vector whose i^{th} entry is the strategy for all players at the i^{th} stage of the sequential game. A non-cooperative sequential game has one well-studied solution: the Subgame Perfect Nash Equilibrium (SPNE).

A strategy profile is an SPNE if its entries are the Nash equilibria of the subgame resulting at each stage of the sequential game. The SPNE of a sequential game captures an equilibrium solution such that no player can make more profit by unilaterally deviating from this strategy in every subgame.

Since the above non-cooperative game is a game with finite number of stages and perfect information, it can be solved using backward induction. Backward induction is the process of reasoning backwards in time, starting from the last stage of the sequential game, to determine a sequence of optimal strategies. It proceeds by first determining the optimal strategies in the last stage. Using this information, one can then decide the optimal strategies for the second-to-last stage of the game. This process continues backwards until the optimal strategies for every stage has been determined. We apply backward induction to the three stage game described above as follows.

Stage 3, Users' decisions: Each consumer located at $x \in [0, 1]$ can choose the services provided by either SP_1 or SP_2 based on its valuation function in (3). The resulting optimal strategy for the consumer is to choose the SP whose index is given by

$$\arg \max_{i \in \{1, 2\}} v_i + t(x - x_i)\varepsilon_i. \quad (5)$$

Since the consumer's utility is a linear function of the normalized privacy risk x and the market is completely covered by the SPs, there exists a threshold x_τ such that the consumer located at x_τ is indifferent to using services provided by SP_1 or SP_2 . Thus, at the indifference threshold x_τ , we have

$$\begin{aligned} u_2(x_\tau) &= u_1(x_\tau) \\ \implies v_2 + t(x_\tau - x_2)\varepsilon_2 &= v_1 + t(x_\tau - x_1)\varepsilon_1. \end{aligned} \quad (6)$$

Simplifying further, the indifference threshold for choosing between the two SPs is given by

$$x_\tau = \frac{v_1 - v_2 + t(F_E(\varepsilon_2)\varepsilon_2 - F_E(\varepsilon_1)\varepsilon_1)}{t(\varepsilon_2 - \varepsilon_1)}, \quad (7)$$

where x_1 and x_2 have been replaced by their corresponding normalized privacy risk values. Thus, given the SPs' tuples $(\varepsilon_i, v_i), i \in \{1, 2\}$, the optimal strategy of a consumer located at x is to use the service of SP_1 if $x \leq x_\tau$ and SP_2 otherwise.

Stage 2, SPs determine QoS: In the second stage, for a given privacy risk guarantee ε_i , SP_i chooses its QoS v_i to maximize its profit π_i . Since a consumer's normalized privacy risk tolerance denotes the fraction of the population

whose privacy risk tolerance is at most ε , x_τ determines the proportion of consumers who choose SP_1 , i.e., n_1 . As a result, the profit functions of SP_1 and SP_2 can be written as

$$\pi_1(v_1; \varepsilon_1; v_2; \varepsilon_2) = [R(\varepsilon_1) - C(v_1; \varepsilon_1)] \frac{v_1 - v_2 + t(F_E(\varepsilon_2)\varepsilon_2 - F_E(\varepsilon_1)\varepsilon_1)}{t(\varepsilon_2 - \varepsilon_1)}, \quad (8)$$

$$\pi_2(v_1; \varepsilon_1; v_2; \varepsilon_2) = [R(\varepsilon_2) - C(v_2; \varepsilon_2)] \left[1 - \frac{v_1 - v_2 + t(F_E(\varepsilon_2)\varepsilon_2 - F_E(\varepsilon_1)\varepsilon_1)}{t(\varepsilon_2 - \varepsilon_1)}\right]. \quad (9)$$

To find the SPNE in this stage, we use the best response method [19]. The best response is a function which captures the behavior of each player while fixing the strategies of the other players. For any $v_{-i} \in \mathcal{V}_{-i}$, we define $BR_i(v_{-i})$ as the best strategy of SP_i such that

$$BR_i(v_{-i}) = \arg \max_{v_i} \pi_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i}), i \in \{1, 2\}. \quad (10)$$

In the Nash equilibrium, each player plays the best response with respect to other players' strategies. Thus, a Nash equilibrium in this stage is a profile $\vec{v}^* = (v_1^*, v_2^*)$ for which

$$v_i^* \in BR_i(v_{-i}), \forall i \in \{1, 2\}. \quad (11)$$

To find the Nash equilibria, we first calculate the best response function of each SP, then find a strategy profile \vec{v}^* for which $v_i^* \in BR_i(v_{-i}), \forall i \in \{1, 2\}$. For a given set of privacy risk guarantees $\{\varepsilon_1, \varepsilon_2\}$, the optimal QoS v_i^* of $SP_i, i \in \{1, 2\}$ in the SPNE is then determined by the solution to the following set of simultaneous equations

$$\begin{aligned} v_1^* &= \arg \max_{v_1} \pi_1(v_1; \varepsilon_1; v_2; \varepsilon_2), \\ v_2^* &= \arg \max_{v_2} \pi_2(v_1; \varepsilon_1; v_2; \varepsilon_2). \end{aligned} \quad (12)$$

Stage 1, SPs determine privacy risk guarantee: In the first stage, we compute equilibrium strategies ε_1 and ε_2 that the two SPs should advertise for optimal market share. Note that x_τ, v_1^* , and v_2^* have been computed in stages 1 and 2 for a fixed ε_1 and ε_2 , and therefore, are functions of ε_1 and ε_2 . The objective functions π_1 and π_2 are thus also functions of ε_1 and ε_2 ; this in turn implies they can be maximized to find the equilibrium strategy ε_1^* and ε_2^* using the best response method.

III. TWO-SP MARKET WITH LINEAR COST AND REVENUE FUNCTIONS

Thus far, we have considered a general model for the consumer distribution of privacy preferences. To obtain better intuition and meaningful analytical solutions, we consider a linear cost and revenue model for each SP. We define the cost function of SP_i to be

$$C(v_i; \varepsilon_i) = cv_i + c\lambda\varepsilon_i, i \in \{1, 2\}, \quad (13)$$

where c and λ are constant scale factors in units of cost/QoS and QoS/privacy risk, respectively. In addition, we model the revenue of each SP from offering a privacy guaranteed service by a linear function

$$R(\varepsilon_i) = r\varepsilon_i + p_i, i \in \{1, 2\}, \quad (14)$$

where r is the revenue per unit privacy risk for using consumers' private data. The parameters p_1 and p_2 model the fixed revenues of the SPs that are independent of consumers' private data.

A. Consumers with Uniformly Distributed Privacy Risk Tolerance

We assume consumers have uniformly distributed privacy risk tolerance between 0 and $\bar{\varepsilon}$. The resulting normalized privacy risk of each SP is given by

$$x_i = F_E(\varepsilon_i) = \frac{\varepsilon_i}{\bar{\varepsilon}}, i \in \{1, 2\}.$$

Let

$$\alpha = \frac{r}{c} - \lambda \quad (15)$$

and

$$\tilde{C} = ct\bar{\varepsilon}. \quad (16)$$

Note that α is the ratio of net profit from using consumer data for a unit of privacy risk to the cost for providing a unit of QoS. Furthermore, \tilde{C} is the cost of providing non-zero utility to the consumer with a maximal mismatch of privacy risk (relative to SP).

By using the backward induction method, the computed SPNE of the two-SP non-cooperative game is presented in the following theorem.

Theorem 1: There exists an SPNE for the two-SP non-cooperative game if the model parameters $\{c, \alpha, t, \bar{\varepsilon}, p_1, p_2\}$ satisfy

$$-1 \leq \frac{16(p_2 - p_1)}{9ct\bar{\varepsilon}} \leq 1, \quad (17)$$

$$\frac{4\alpha - 3t}{3t} \leq \frac{16(p_2 - p_1)}{9ct\bar{\varepsilon}} \leq \frac{4\alpha - t}{3t}, \quad (18)$$

$$(12c\alpha\bar{\varepsilon})^2 - (15ct\bar{\varepsilon})^2 + 288ct\bar{\varepsilon}(p_2 + p_1) \geq [16(p_2 - p_1)]^2. \quad (19)$$

The closed form solution of the SPNE is given by

$$\varepsilon_2^* = \frac{12\bar{\varepsilon}c\alpha + 15ct\bar{\varepsilon} - 16(p_2 - p_1)}{24tc}, \quad (20)$$

$$v_2^* = \frac{(2\alpha + t)c\alpha 6\bar{\varepsilon} + (\alpha - t)9ct\bar{\varepsilon} + (t - 2\alpha)8p_2 + (\alpha + t)16p_1}{24ct}, \quad (21)$$

$$\varepsilon_1^* = \varepsilon_2^* - \frac{3\bar{\varepsilon}}{4}, \quad (22)$$

$$v_1^* = v_2^* - \frac{3\bar{\varepsilon}}{4}\alpha + \frac{p_2 - p_1}{3c}. \quad (23)$$

At the equilibrium, i.e., for the SPNE, $\frac{1}{2} - \frac{8(p_2 - p_1)}{9ct\bar{\varepsilon}}$ of the population choose the service provided by SP_1 while the remaining $\frac{1}{2} + \frac{8(p_2 - p_1)}{9ct\bar{\varepsilon}}$ of the population choose SP_2 . The total profit of SP_1 and SP_2 are given by

$$\pi_1^* = \frac{4c}{27t\bar{\varepsilon}} \left(\frac{9t\bar{\varepsilon}}{8} - \frac{2(p_2 - p_1)}{c} \right)^2; \pi_2^* = \frac{4c}{27t\bar{\varepsilon}} \left(\frac{9t\bar{\varepsilon}}{8} + \frac{2(p_2 - p_1)}{c} \right)^2.$$

The proof of Theorem 1 is provided in Appendix A.

In theorem 1, we observe $\varepsilon_2 = \varepsilon_1 + \frac{3}{4}\bar{\varepsilon}$, which implies for any SPNE strategy profile, the SP who with the higher QoS will offer a privacy risk guarantee $\frac{3}{4}\bar{\varepsilon}$ more than its competitor. Furthermore, a higher maximum net profit for using private data, denoted by $c\alpha\bar{\varepsilon}$, encourages SPs to offer higher privacy risk guarantees. However, a larger difference in revenues independent of private information between SP_1 and SP_2 , denoted by $p_2 - p_1$, encourages SPs to lower their privacy risk guarantees to attract more consumers with lower privacy risk tolerance. Also, a larger \tilde{C} will result in a larger differentiation in privacy risk guarantee in the equilibrium strategies of SP_1 and SP_2 . Finally, as $p_2 - p_1$ increases, the market share of SP_1 decreases while SP_2 's market share increases.

By (16), x_τ , π_1^* , and π_2^* can be simplified to

$$x_\tau = \frac{1}{2} - \frac{8(p_2 - p_1)}{9\tilde{C}}. \quad (24)$$

$$\pi_2^* = \frac{1}{3} \left(\frac{3}{4} \sqrt{\tilde{C}} + \frac{4(p_2 - p_1)}{3\sqrt{\tilde{C}}} \right)^2 \quad \pi_1^* = \frac{1}{3} \left(\frac{3}{4} \sqrt{\tilde{C}} - \frac{4(p_2 - p_1)}{3\sqrt{\tilde{C}}} \right)^2 \quad (25)$$

Note that for a fixed $p_2 - p_1$, both π_1^* and π_2^* are decreasing functions of \tilde{C} when $\tilde{C} \in [0, \frac{16(p_2 - p_1)}{9}]$ and increasing afterwards. On the other hand, by (17), we have $\frac{16(p_2 - p_1)}{9} \leq c\alpha\bar{\varepsilon}$, which implies $\tilde{C} \geq \frac{16(p_2 - p_1)}{9}$. Therefore, both π_1^* and π_2^* are increasing function of \tilde{C} in the SPNE. This indicates both SPs will make more profits in the SPNE with a larger \tilde{C} .

B. Consumers with Normally Distributed Privacy Risk Tolerance

In this section, we model consumers' privacy tolerance as a random variable E that follows a normal distribution $\mathcal{N}(\frac{\bar{\varepsilon}}{2}, \sigma^2)$ with a mean of $\frac{\bar{\varepsilon}}{2}$ and a standard deviation of σ . Since $E \in [0, \bar{\varepsilon}]$, we restricted the normal distribution to lie within the interval $[0, \bar{\varepsilon}]$. Then E conditional on $E \in [0, \bar{\varepsilon}]$ follows a truncated normal distribution with cumulative distribution function

$$F_E(\varepsilon) = \begin{cases} \frac{\Phi(\frac{\varepsilon - \frac{\bar{\varepsilon}}{2}}{\sigma}) - \Phi(-\frac{\frac{\bar{\varepsilon}}{2}}{\sigma})}{\Phi(\frac{\bar{\varepsilon}}{2\sigma}) - \Phi(-\frac{\bar{\varepsilon}}{2\sigma})} & \varepsilon \in [0, \bar{\varepsilon}] \\ 0 & \varepsilon \in [-\infty, 0] \\ 1 & \varepsilon \in [\bar{\varepsilon}, +\infty] \end{cases}, \quad (26)$$

where $\Phi(y)$ denotes the CDF of the standard normal distribution.

In contrast to the uniform distribution case, the CDF in (26) is not amenable to a closed form solution. Thus, we characterize the equilibrium numerically. To find the SPNE, we first compute the SPNE QoS in the second stage as functions of privacy risk guarantees by solving (12). Then, we use an iterated best response method to find the optimal privacy risk guarantee of an SP by fixing other SPs strategies in each iteration. When the process converges, we have found an SPNE in which no SP is better off by unilaterally deviating from the equilibrium.

C. Illustration of Results

In this section, we illustrate our model and results. First, we assume consumers have uniformly distributed privacy risk tolerance. We plot each SP's SPNE strategies, market share, and total profit with respect to consumers'

maximum privacy risk tolerance for different values of the QoS/privacy risk scale factor t . Later, we study the model in which consumers' privacy risk tolerance follows a normal distribution $\mathcal{N}(\frac{\bar{\varepsilon}}{2}, 1)$ truncated between 0 and $\bar{\varepsilon}$. The model parameters are given as follows:

Parameter	c	λ	r	p_1	p_2
Value	0.5	0.75	0.7	0.4	0.8

Table. I: Numerical Example Model Parameters

1) *Consumers with Uniformly Distributed Privacy Risk Tolerance:* In this section, we vary $\bar{\varepsilon}$ from 3 to 5 to study properties of SPNE. Note that by theorem 1, t must belong to $[0.58, 0.85]$ for a stable and sustainable SPNE. In Figure 2, we plot the equilibrium strategies of different SPs. Observe that both privacy risks and QoSs are linear functions of $\bar{\varepsilon}$, as expected for the linear model assumed. Furthermore, it can be seen that as t , the valuation of privacy by consumer, decreases, each SP will increase its privacy risk to generate more profit from using private data. Correspondingly, the SPs will have to provide higher QoS to attract consumers.

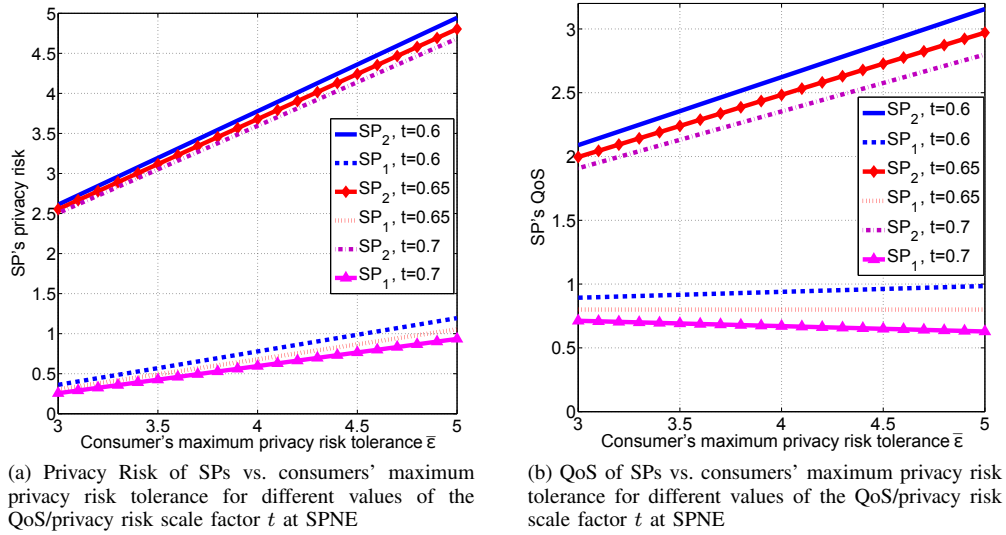


Figure 2: Equilibrium strategies of SPs vs. consumers' maximum privacy risk tolerance for different values of the QoS/privacy risk scale factor t under a uniform distribution of consumer privacy risk

The market shares of different SPs in the SPNE are presented in Figure 3. We observe that the equilibrium market share of SP_2 decreases as t increases. The intuition behind this is that if t increases, the consumer's valuation of privacy mismatch also increases. Thus, it is more difficult for SP_2 to attract consumers with privacy tolerance lower than ε_2 . As a result, its market share decreases. Notice that in Figure 3, as $\bar{\varepsilon}$ decreases, the equilibrium market share of SP_2 also increases. This is because consumers experience a smaller negative utility from the mismatch between their preferred and the offered privacy risk when the net range is smaller (recall that the utility from mismatch is given by $t(x - x_i)\varepsilon_i, \varepsilon_i \in [0, \bar{\varepsilon}]$). Thus, SP_2 has the incentive to offer high QoS with high privacy risk. As a result, more consumers will choose the SP with a higher privacy risk to enjoy a higher QoS.






Figure 3: Market shares of SPs at SPNE vs. consumers' maximum privacy risk tolerance for different values of the QoS/privacy risk scale factor t under a uniform distribution of consumer privacy risk

In Figure 4, we plot the total profit at the SPNE for each SP as a function of the maximum consumer privacy risk tolerance $\bar{\epsilon}$ for different values of t . As shown in the figure, the total profit of both SPs at SPNE increases as $\bar{\epsilon}$ increases. This is due to the fact that a larger $\bar{\epsilon}$ indicates a larger range of consumer preferences, and then, more possibilities for the SPs to exploit private information. Thus, both SPs can benefit from the using the private data of consumers that have a higher privacy risk tolerance. As t decreases, the total profit of both SPs decrease. This is due to the fact that as t decreases, the SPs intend to increase their privacy risks to generate more profit. As a result, the QoS of each SP increases to attract more consumers (see Figure 2), that in turn has the consequence of increasing the cost of providing services. Thus, the SPs have lower profits.



(a) Total profit of SP_2 vs. consumers' maximum privacy risk tolerance for different values of the QoS/privacy risk scale factor t at SPNE



(b) Total profit of SP_1 vs. consumers' maximum privacy risk tolerance for different values of the QoS/privacy risk scale factor t at SPNE

Figure 4: Total equilibrium profit of SPs vs. consumers' maximum privacy risk tolerance for different values of the QoS/privacy risk scale factor t under a uniform distribution of consumer privacy risk

2) *Consumers with Normally Distributed Privacy Risk Tolerance:* We now consider the case in which consumers' privacy risk tolerance follows a truncated normal distribution with a mean of $\bar{\varepsilon}$ and a standard deviation of 1. The equilibrium strategies of different SPs are shown in Figure 5. As with the uniform distribution scenario, here too we observe that the privacy risk and the QoS offered by each SP are linear functions of $\bar{\varepsilon}$. We also notice that in the SPNE, SP_2 will always provide service with maximum privacy risk (Figure 5a). This is because the truncated normal distribution forces SP_2 to forfeit privacy differentiation while maximizing its profit. This allows SP_1 to gain an advantage (relative to the uniform distribution). Furthermore, we observe that the value of t has slightly less influence on strategies of SPs in the SPNE compared to uniformly distributed consumer privacy tolerance.




Figure 5: Equilibrium strategies of SPs vs. consumers' maximum privacy risk tolerance for different values of t under truncated normal privacy risk tolerance distribution

Figure 6 shows market shares of different SPs at SPNE vs. consumers' maximum privacy risk tolerance for different values of t under truncated normal privacy tolerance distribution. As t decreases, the market share of SP_2 at SPNE increases, and vice versa. Also, when $\bar{\varepsilon}$ decreases, the equilibrium market share of SP_1 also decreases. Furthermore, it can be seen that for the same $\bar{\varepsilon}$, the market share of SP_2 (SP_1) is smaller (larger) when consumers' privacy tolerance follows the truncated normal distribution compared to uniform distribution. Furthermore, our numerical analysis shows that at SPNE, SP_2 is forced to provide service with maximum privacy risk. We argue that this is due to the shape of the distribution that limits the number of consumers at the two extremes thus compelling the two SPs compete for the large bulk of consumers distributed around $\bar{\varepsilon}/2$. Given the ability of SP_2 to make more profit on untargeted services relative to SP_1 , the SPNE solution leads to SP_1 increasing its market share to be profitable and SP_2 achieving profitability with a smaller market share.

The relationship between total profit of different SPs at SPNE vs. consumers' maximum privacy risk tolerance for different values of t is shown in Figure 7. Similar to Figure 4, both SPs' total profit increases as $\bar{\varepsilon}$ increases. However, in contrast to Figure 4, as t decreases, the total profit of SP_2 increases. This is because SP_2 always offers $\bar{\varepsilon}$ in the SPNE. Notice that SP_2 's equilibrium QoS is also a linear function of ε (see Figure 5b). On the




Figure 6: Market shares of SPs at SPNE vs. consumers' maximum privacy risk tolerance for different values of t under truncated normal privacy tolerance distribution

other hand, SP_2 's market share increases as t decreases (see Figure 6). By (4), (13), and (14); the total profit of SP_2 increases as t decreases.

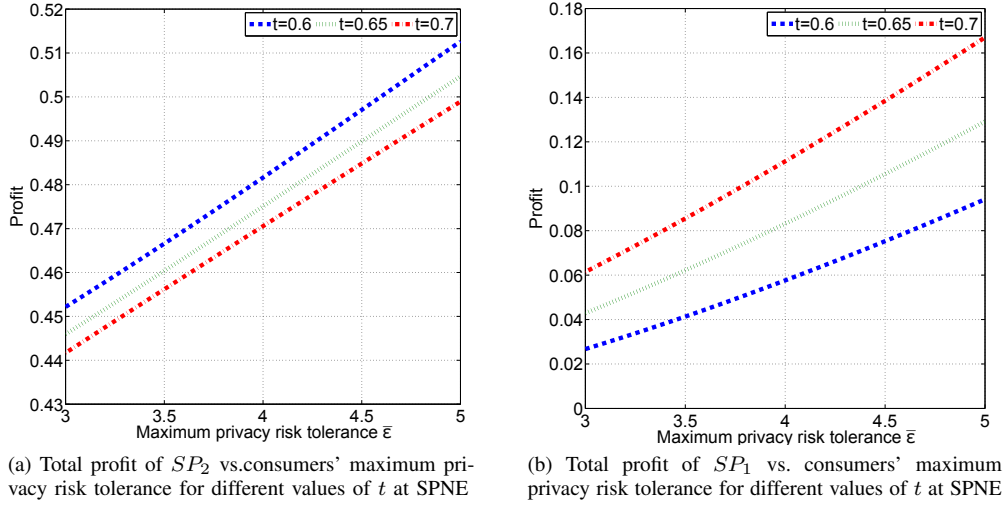


Figure 7: Total equilibrium profit of SPs vs. consumers' maximum privacy risk tolerance for different values of t under truncated normal privacy risk tolerance distribution

IV. MARKET WITH MULTIPLE SERVICE PROVIDERS

In the previous sections, we studied the market with two SPs. In this section, we examine a generalized model with multiple SPs (Figure 8). We allow for a finitely arbitrary number of SPs, each of which offers the same type of free service but with different QoS and privacy risk guarantee to consumers. In particular, we assume there are m SPs in the market. Our models for cost, revenue and utility for each SP as well as the consumers are the same as for the two-SP model described in Section II-A3. Furthermore, we assume a consumer's privacy risk tolerance is uniformly distributed between $[0, \varepsilon]$. Analogous to the two SP model, the interactions between the m SPs and consumers can also be viewed as a non-cooperative sequential game. The m -SP game proceeds in three stages.

In the first stage, each of the m SPs chooses its own privacy risk guarantee resulting in a vector $\vec{\varepsilon} = (\varepsilon_1, \varepsilon_2, \dots, \varepsilon_m)$ (on the interval $[0, \bar{\varepsilon}]$). Without loss of generality, we assume $\varepsilon_1 \leq \varepsilon_2 \leq \dots \leq \varepsilon_m$. At the second stage, given the privacy risk $\vec{\varepsilon}$ determined in the first stage, the SPs simultaneously determine their QoS values to obtain a vector $\vec{v} = \{v_1, v_2, \dots, v_m\}$. At the last stage, each consumer chooses the SP that yields the maximal perceived utility for the consumer.




Figure 8: Market model for multiple SPs offering services with privacy guarantee

To find the SPNE, we apply backward induction to the three stage game described above as follows. In the last stage of the game, for fixed QoS and privacy risk guarantee strategies of the SPs, consumers' choices of SPs are determined by their privacy risk tolerances. In the two-SP case, the consumer located at x_τ divides the set of consumers into two convex subsets where the consumers in the left subset will choose SP_1 and vice versa. However, for the multiple SP case, the market share of SP_i ($i \in \{1, \dots, m\}$) is not necessarily a convex set between the indifference threshold in which consumers are indifferent to choosing SP_{i-1} or SP_i and the threshold in which consumers are indifferent to choosing SP_i or SP_{i+1} . This is due to the fact that in general the problem requires each SP_i to compete with all other SPs, even if their privacy risk offerings are very different (e.g., SPs with a large difference in locations in Figure 8). We note that this will not happen in the equilibrium since an SP with zero market share would be better off by either improving its QoS to attract some consumers or just exit the market. Therefore, in the equilibrium, each SP only competes directly with its two closest neighbors. For given QoS profile $\vec{v} = \{v_1, v_2, \dots, v_m\}$ and privacy risk profile $\vec{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_m)$, the market share of each SP are

$$\begin{aligned}
 n_1 &= \frac{v_1 - v_2 + t(F_E(\varepsilon_2)\varepsilon_2 - F_E(\varepsilon_1)\varepsilon_1)}{t(\varepsilon_2 - \varepsilon_1)}, \\
 n_i &= \frac{v_i - v_{i+1} + t(F_E(\varepsilon_{i+1})\varepsilon_{i+1} - F_E(\varepsilon_i)\varepsilon_i)}{t(\varepsilon_{i+1} - \varepsilon_i)} - \frac{v_{i-1} - v_i + t(F_E(\varepsilon_i)\varepsilon_i - F_E(\varepsilon_{i-1})\varepsilon_{i-1})}{t(\varepsilon_i - \varepsilon_{i-1})}, \\
 i &\in \{2, \dots, m-1\}, \\
 n_m &= 1 - \frac{v_{m-1} - v_m + t(F_E(\varepsilon_m)\varepsilon_m - F_E(\varepsilon_{m-1})\varepsilon_{m-1})}{t(\varepsilon_m - \varepsilon_{m-1})}.
 \end{aligned}$$

Furthermore, we define the objective functions of SP_i to be

$$\pi_i(\vec{\varepsilon}; \vec{v}) = [R(\varepsilon_i) - C(v_i; \varepsilon_i)]n_i(\vec{\varepsilon}; \vec{v}), i \in \{1, \dots, m\}.$$

For a given privacy risk guarantees profile $\vec{\varepsilon}$, the optimal QoS of SP_i ($i \in \{1, \dots, m\}$) is determined by

$$\arg \max_{v_i} \pi_i(\vec{\varepsilon}; \vec{v}), i \in \{1, \dots, m\} \quad (27)$$

while fixing all other players' strategies.

We note that the cost function $C(v_i; \varepsilon_i)$ and the market segmentation computed in the first stage are both linear functions of v_i . Thus, for a fixed privacy risk guarantee profile $\vec{\varepsilon}$, the objective functions of SP_i in this stage is a concave function with respect to its own strategy v_i . Furthermore, the feasible set of each SP's strategy is a convex set. Thus, the non-cooperative game among the SPs in this stage can be a m-player concave game. By [20], there exists a Nash equilibrium. We define $\delta_i \triangleq \frac{1}{t(\varepsilon_{i+1} - \varepsilon_i)}$, $y_1 \triangleq r(\varepsilon_1) + p_1 - c\lambda\varepsilon_1 - ct x_2 \varepsilon_2 + ct x_1 \varepsilon_1$, $y_N \triangleq r(\varepsilon_N) + p_N - c\lambda\varepsilon_N - ct(1 - x_N)\varepsilon_N + ct(1 - x_{N-1})\varepsilon_{N-1}$ and $y_i \triangleq \frac{r(\varepsilon_i) + p_i - c\lambda\varepsilon_i + ct x_i \varepsilon_i - ct x_{i+1} \varepsilon_{i+1}}{t(\varepsilon_{i+1} - \varepsilon_i)} + \frac{r(\varepsilon_i) + p_i - c\lambda\varepsilon_i - ct x_{i-1} \varepsilon_{i-1} + ct x_i \varepsilon_i}{t(\varepsilon_i - \varepsilon_{i-1})} \quad \forall i \in \{2, \dots, m\}$. Applying the first order condition to SPs' profit functions (solving simultaneous linear equations obtained from $\frac{\partial \pi_i(\vec{\varepsilon}; \vec{v})}{\partial v_i} = 0, i \in \{1, 2, \dots, m\}$) yields the equilibrium strategies

$$v_1^* = \frac{v_2}{2} + \frac{y_1}{2c}, \quad (28)$$

$$v_i^* = \frac{cv_{i+1}\delta_i + cv_{i-1}\delta_{i-1} + y_i}{2c[\delta_i + \delta_{i-1}]}, i \in \{2, \dots, m\}, \quad (29)$$

$$v_m^* = \frac{v_{m-1}}{2} + \frac{y_m}{2c}. \quad (30)$$

In the last stage, the SPs determine their privacy risk guarantees $\vec{\varepsilon}$ by considering equilibrium strategies in previous stages (n_i and $v_i^* \quad \forall i \in \{1, \dots, m\}$) as functions of $\vec{\varepsilon}$. Therefore, the optimal privacy risk strategy of SP_i is determined by

$$\arg \max_{\varepsilon_i} \pi_i(\vec{\varepsilon}; \vec{v}), i \in \{1, \dots, m\} \quad (31)$$

while fixing all other players' strategies. For reasons of intractability (solving highly parameterized high order polynomial equations), a full characterization of privacy risk equilibria could not be achieved. Thus, we characterize the SPNE numerically by using the iterated best response method. We consider a three-SP market and adopt the model parameters presented in Table I. Furthermore, we assume $t = 0.7$ and $\bar{\varepsilon} = 5$. The initial privacy risk of SP_i is given by $\frac{i\bar{\varepsilon}}{i+1}$ for $i \in \{1, 2, 3\}$. Although there exists an SPNE in the second stage of the sequential game for fixed privacy guarantees, the existence of an equilibrium in the first stage can not be guaranteed.

The best response strategies of the SPs for different values of SP_2 's privacy independent revenue are plotted in Figure 9. It can be seen that the two SPs with lower privacy risks proceed to jump over each other in each round of best response iteration, attempting to lower its privacy risk to attract more consumers from its competitor. The SP with the highest revenue independent of using private data adopts a high privacy risk strategy to focus on consumers with high privacy risk tolerance and exploiting their private data extensively. Furthermore, we observe that when p_2 is large, SP_3 's privacy risk strategy is also higher on average. On the other hand, SP_1 's best response strategy is lower. The intuition behind is that a larger p_2 allows SP_2 to set a higher privacy risk to make more profit from using consumer data. This forces SP_3 to increase its privacy risk to differentiate itself from SP_2 . On the other




Figure 9: Best response of each SP's privacy risk for different values of SP_2 's revenue independent of using private data

hand, a higher privacy risk of SP_2 will encourage SP_1 to lower its privacy risk to attract more consumers.

V. CONCLUSIONS

Our work seeks to understand the effect of offering privacy- and QoS- differentiated online services on consumers with heterogeneous expressions of privacy sensitivity. We have quantified this effect as the fraction of consumers that prefer lower privacy risks with the accompanying lower QoS to the alternative of higher risks and higher QoS. We have presented an analysis built upon the classical Hotelling model to compute these fractions for both the two and multi SP problem. Analogous to the classical segmentation models, our problem also involves parameters that capture cost, revenue, and consumer valuation functions that are dependent and independent of the privacy risks. While such parametrized model can make the analysis challenging, our results for relatively simple yet meaningful functions such as linear cost models and uniform (as well as truncated Gaussian) distribution of consumer preferences suggests that SPs that have higher profits from untargeted services have an edge in the market. SPs competing on offering lower privacy risks have to offer better QoS or figure out other means of increasing untargeted revenue to gain market share. The market segmentation model assumes at least two or more SPs were able to overcome the barrier to entry and differentiate themselves. Thus, a related question we will address going forward is whether such barriers to entry are in fact surmountable when competitors use privacy as a differentiator. These analyses are crucial for developing better privacy policies to effectively enable safe and secure online commerce.

VI. ACKNOWLEDGEMENTS

We wish to thank Professor Anand Sarwate at Rutgers University for many interesting discussions. This work is funded in part by the National Science Foundation under grant CCF-1350914.

APPENDIX

APPENDIX

Starting from the last stage in which consumers choose different SPs, we use backward induction to find the SPNE of the sequential game. In the last stage, each consumer located at $x \in [0, 1]$ chooses an SP which maximize

its utility function (3). By (7) and the assumption that consumers' privacy risk tolerances are uniformly distributed, the indifference threshold x_τ is given by

$$x_\tau = \frac{v_1 - v_2 + \frac{t(\varepsilon_2^2 - \varepsilon_1^2)}{\bar{\varepsilon}}}{t(\varepsilon_2 - \varepsilon_1)} = n_1(v_1; \varepsilon_1; v_2; \varepsilon_2). \quad (32)$$

At the second stage, the optimal strategy of each SP is determined by the solution of (12). For fixed privacy risk guarantees ε_2 and ε_1 , the objective function of $SP_i, i \in \{1, 2\}$ in this stage, i.e. $\pi_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i})$, is a concave function with respect to its own strategy v_i . Furthermore, the feasible set of SP_i 's strategy is a convex set ($v_i \in [0, +\infty]$). Thus, the non-cooperative subgame between SP_2 and SP_1 in this stage can be considered as a two-player concave game. By Theorem 1 and 2 in [20], we can establish

Proposition 1: For fixed privacy risk strategies, there exists a unique Nash equilibrium in the game between SP_2 and SP_1 at the second stage.

To compute the equilibrium strategy of the second stage, we first substitute (13), (14), and (32) into (9) and (8). Then, we apply the first order condition to SPs' profit functions and solve the simultaneous equations given by

$$\frac{\partial \pi_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i})}{\partial v_i} = 0 \quad \forall i \in \{1, 2\}. \quad (33)$$

Solving the above simultaneous equations yields

$$v_1 = \frac{r\varepsilon_1 + p_1}{2c} + \frac{v_2 - \lambda\varepsilon_1 - tx_2\varepsilon_2 + tx_1\varepsilon_1}{2}, \quad (34)$$

$$v_2 = \frac{r\varepsilon_2 + p_2}{2c} + \frac{v_1 - \lambda\varepsilon_2 - t(1 - x_2)\varepsilon_2 + t(1 - x_1)\varepsilon_1}{2}. \quad (35)$$

For given privacy guarantees ε_1 , and ε_2 , solving the simultaneous linear equations above by substituting (34) into (35) yields the equilibrium strategies

$$v_1^*(\varepsilon_2, \varepsilon_1) = \frac{2(r\varepsilon_1 + p_1) + r\varepsilon_2 + p_2}{3c} + \frac{t(1 + x_1)\varepsilon_1 - \lambda(\varepsilon_2 + 2\varepsilon_1) - t(1 + x_2)\varepsilon_2}{3}, \quad (36)$$

$$v_2^*(\varepsilon_2, \varepsilon_1) = \frac{2(r\varepsilon_2 + p_2) + r\varepsilon_1 + p_1}{3c} + \frac{t(2 - x_1)\varepsilon_1 - \lambda(2\varepsilon_2 + \varepsilon_1) - t(2 - x_2)\varepsilon_2}{3}. \quad (37)$$

At the first stage, the SPs determine their optimal privacy risk by considering the QoS of each SP and the market segmentation computed in previous stages as functions of privacy risks offered by the SPs. By substituting (37) and (36) into (9) and (8), the profit functions of the SPs can be written as

$$\pi_2 = \frac{c}{9t(\varepsilon_2 - \varepsilon_1)} \left[\frac{p_2 - p_1}{c} + \left(\frac{r}{c} - \lambda + t \frac{2\bar{\varepsilon} - \varepsilon_2 - \varepsilon_1}{\bar{\varepsilon}} \right) (\varepsilon_2 - \varepsilon_1) \right]^2, \quad (38)$$

$$\pi_1 = \frac{c}{9t(\varepsilon_2 - \varepsilon_1)} \left[-\frac{p_2 - p_1}{c} + \left(-\frac{r}{c} + \lambda + t \frac{\bar{\varepsilon} + \varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}} \right) (\varepsilon_2 - \varepsilon_1) \right]^2. \quad (39)$$

Next, we apply the first order condition to SPs' profit functions to compute the equilibrium strategies. Taking the

derivatives of π_2 and π_1 with respect to ε_2 and ε_1 and set both of their values to 0 yields

$$\frac{\partial \pi_2}{\partial \varepsilon_2} = \frac{c[(\frac{r}{c} - \lambda + t \frac{2\bar{\varepsilon} - \varepsilon_2 - \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) + \frac{p_2 - p_1}{c}][(\frac{r}{c} - \lambda + t \frac{2\bar{\varepsilon} - 3\varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c}]}{9t(\varepsilon_2 - \varepsilon_1)^2} = 0, \quad (40)$$

$$\frac{\partial \pi_1}{\partial \varepsilon_1} = \frac{c[(-\frac{r}{c} + \lambda + t \frac{\bar{\varepsilon} + \varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c}][(\frac{r}{c} - \lambda - t \frac{\bar{\varepsilon} - \varepsilon_2 + 3\varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c}]}{9t(\varepsilon_2 - \varepsilon_1)^2} = 0. \quad (41)$$

Solving the two simultaneous equations above yields

$$(\frac{r}{c} - \lambda + t \frac{2\bar{\varepsilon} - \varepsilon_2 - \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) + \frac{p_2 - p_1}{c} = 0 \quad (42)$$

or

$$(\frac{r}{c} - \lambda + t \frac{2\bar{\varepsilon} - 3\varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c} = 0 \quad (43)$$

and

$$(-\frac{r}{c} + \lambda + t \frac{\bar{\varepsilon} + \varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c} = 0 \quad (44)$$

or

$$(\frac{r}{c} - \lambda - t \frac{\bar{\varepsilon} - \varepsilon_2 + 3\varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c} = 0. \quad (45)$$

We note that the strategies given by (42) and (44) result in 0 profits in (38) and (39). This indicates the privacy risk determined by (42) and (44) are strictly dominated by the strategies given by the solution of (43) and (45). Solving (43) and (45) yields the equilibrium privacy risk (20) and

$$\varepsilon_1^* = \frac{12\bar{\varepsilon}c\alpha - 3ct\bar{\varepsilon} - 16(p_2 - p_1)}{24tc}. \quad (46)$$

By subtracting (46) from (20), we have (22). Substitute the solution of ε_2^* and ε_1^* to (37) and (36), we have (21) and

$$v_1^* = \frac{(2\alpha - t)c\alpha 6\bar{\varepsilon} + (\alpha - 3t)3ct\bar{\varepsilon} + (t - \alpha)16p_2 + (2\alpha + t)8p_1}{24ct}. \quad (47)$$

Subtracting (47) from (21) yields (23).

Next, we prove the sufficient condition for the existence of the above SPNE. First of all, the model parameters must sustain a competitive market environment. Thus, in the equilibrium, each SP must have non-zero market share. This indicates the parameters must satisfy $0 \leq x_\tau^* = \frac{v_1^* - v_2^* + t(x_2^* \varepsilon_2^* - x_1^* \varepsilon_1^*)}{t(\varepsilon_2^* - \varepsilon_1^*)} \leq 1$. Substitute (20), (21), (22), and (23) into the above inequality, we have (17). Furthermore, in the SPNE, the QoS of each SP must be non-negative (QoS feasibility) and the privacy risk guarantees must be bounded between 0 and $\bar{\varepsilon}$ (privacy risk feasibility). By the model assumption in Section II-A1, we have $\varepsilon_1 \leq \varepsilon_2$. Thus, we only requires $\varepsilon_2 \leq \bar{\varepsilon}$ and $\varepsilon_1 \geq 0$. Substitute (20) and (22) into the two inequalities above yields (18). Let $x_i^* = \frac{\varepsilon_i^*}{\bar{\varepsilon}}, i \in \{A, B\}$ denote the normalized privacy risk of each SP in the SPNE. The equilibrium strategies must satisfy the complete market coverage condition given by $u_i(x) = v_i^* - t(x - x_i^*)\varepsilon_i^* \geq 0 \quad \forall x \in [0, 1]$ for at least one $i \in \{A, B\}$.

Substituting (18) into (23), we have $v_2^* - v_1^* = \frac{3\bar{\varepsilon}}{4}\alpha - \frac{p_2 - p_1}{3c} \geq \frac{3t\bar{\varepsilon}}{16} + \frac{2(p_2 - p_1)}{3c} > 0$, thus we only need $v_1 \geq 0$ for QoS feasibility. Furthermore, the Hotelling model feasibility condition implies $v_1^* - tx_1^*\varepsilon_1^* \geq v_2^* - tx_2^*\varepsilon_2^*$. Since $u_i(x)$ is an increasing function of x , complete market coverage condition can be simplified to $u_1(0) \geq 0$. As a result, the QoS feasibility condition and the complete market coverage condition can be simplified to $v_1^* - tx_1^*\varepsilon_1^* \geq 0$. Therefore, the sufficient condition for the existence of SPNE is given by:

- 1) $0 \leq \frac{v_1^* - v_2^* + t(x_2^*\varepsilon_2^* - x_1^*\varepsilon_1^*)}{t(\varepsilon_2^* - \varepsilon_1^*)} \leq 1$,
- 2) $0 \leq \varepsilon_1^*, \varepsilon_2^* \leq \bar{\varepsilon}$,
- 3) $v_1^* - tx_1^*\varepsilon_1^* \geq 0$.

Solving the above three inequalities yield (17), (18), and (19). The equilibrium market share and profits of the SPs are obtained by substituting (20), (21), (22), and (23) into (7), (8), and (9).

REFERENCES

- [1] Úlfar Erlingsson, V. Pihur, and A. Korolova, "RAPPOR: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS '14)*, 2014, pp. 1054–1067. [Online]. Available: <http://dx.doi.org/10.1145/2660267.2660348>
- [2] "About privacy and location services in ios 8 and later," <https://support.apple.com/en-is/HT203033>, Sep 2016.
- [3] A. Acquisti, C. Taylor, and L. Wagman, "The economics of privacy," *Journal of Economic Literature*, vol. 54, no. 2, pp. 442–492, jun 2016. [Online]. Available: <http://dx.doi.org/10.1257/jel.54.2.442>
- [4] G. Shaffer and Z. J. Zhang, "Competitive coupon targeting," *Marketing Science*, vol. 14, no. 4, pp. 395–416, 1995.
- [5] Y. Chen and G. Iyer, "Research note consumer addressability and customized pricing," *Marketing Science*, vol. 21, no. 2, pp. 197–208, 2002.
- [6] Z. Tang, Y. Hu, and M. D. Smith, "Gaining trust through online privacy protection: Self-regulation, mandatory standards, or caveat emptor," *Journal of Management Information Systems*, vol. 24, no. 4, pp. 153–173, 2008.
- [7] J. Campbell, A. Goldfarb, and C. Tucker, "Privacy regulation and market structure," *Journal of Economics & Management Strategy*, vol. 24, no. 1, pp. 47–73, 2015.
- [8] V. Conitzer, C. R. Taylor, and L. Wagman, "Hide and seek: Costly consumer privacy in a market with repeat purchases," *Marketing Science*, vol. 31, no. 2, pp. 277–292, 2012.
- [9] Y. Chen, C. Narasimhan, and Z. J. Zhang, "Individual marketing with imperfect targetability," *Marketing Science*, vol. 20, no. 1, pp. 23–41, 2001.
- [10] N. Jentzsch, S. Preibusch, and A. Harasser, "Study on monetising privacy: An economic model for pricing personal information," *ENISA*, Feb, 2012.
- [11] D.-J. Lee, J.-H. Ahn, and Y. Bang, "Managing consumer privacy concerns in personalization: a strategic analysis of privacy protection," *Management Information Systems Quarterly*, vol. 35, no. 2, pp. 423–444, 2011.
- [12] R. K. Chellappa and S. Shivendu, "Mechanism design for "free" but "no free disposal" services: The economics of personalization under privacy concerns," *Management Science*, vol. 56, no. 10, pp. 1766–1780, 2010. [Online]. Available: <http://dx.doi.org/10.1287/mnsc.1100.1210>
- [13] A. Datta, M. C. Tschantz, and A. Datta, "Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination," in *Proceedings on Privacy Enhancing Technologies*, vol. 2015, no. 1, apr 2015, pp. 92–112. [Online]. Available: <http://dx.doi.org/10.1515/popets-2015-0007>
- [14] J. M. Perloff, *Microeconomics: theory and applications with calculus*. Pearson, 2016.
- [15] H. Hotelling, "Stability in competition," in *The Collected Economics Articles of Harold Hotelling*. Springer, 1990, pp. 50–63.
- [16] X. Wauthy, "Quality choice in models of vertical differentiation," *The Journal of Industrial Economics*, pp. 345–353, 1996.
- [17] S. Brenner, "Hotelling games with three, four, and more players," *Journal of Regional Science*, vol. 45, no. 4, pp. 851–864, 2005.
- [18] D. Fudenberg and J. Tirole, *Game theory*. MIT press, MA, 1991.
- [19] M. J. Osborne and A. Rubinstein, *A course in game theory*. MIT press, 1994.

- [20] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica: Journal of the Econometric Society*, pp. 520–534, 1965.